Orthophoto generation with gaussian splatting: mitigating reflective surface artifacts in imagery from low-cost sensors

Mattia Previtali [1], Luigi Barazzetti [1], Fabio Roncoroni [2]

[1] Dept. of Architecture, Built environment and Construction engineering (ABC)
Politecnico di Milano, Via Ponzio 31, Milan, Italy
(mattia.previtali, luigi.barazzetti,)@polimi.it
[2] Polo territoriale di Lecco, via Previati 1/c, Lecco, 23900, Italy - fabio.roncoroni@polimi.it

Keywords: Orthophoto, Gaussian Splatting, Reflective surfaces, Glass, Metal.

Abstract

Orthophotos offer a distortion-free and high-resolution depiction of architectural and mechanical elements, serving critical functions in reverse engineering, defect detection, and condition assessment. Recent advancements in low-cost sensors have made them increasingly popular for orthophoto generation due to their affordability and ease of use. However, reflective surfaces pose significant challenges in traditional photogrammetric workflows, leading to inaccuracies in feature matching and 3D reconstruction. This paper investigates the integration of Gaussian splatting into the orthophoto generation process as a solution to address these challenges. Gaussian splatting is particularly effective in handling irregular and sparse data, making it suitable for scenarios involving reflective surfaces. In this study, datasets containing reflective surfaces, such as metallic elements and urban environments, were used to evaluate the performance of Gaussian splatting compared to traditional photogrammetry workflow. Our findings indicate that Gaussian splatting effectively reduces artifacts caused by reflections while preserving geometric accuracy and critical detail in non-reflective areas. Additionally, this approach proves to be computationally efficient, making it ideal for low-cost sensor applications. Although limitations remain, such as smoothing effects that may reduce fine detail, the proposed methodology shows promise for improving orthophoto quality in complex environments.

1. Introduction

Orthophotos offer a high-resolution, distortion-free view of the current state of a building, site, or mechanical element. They serve critical functions in various applications, including reverse engineering, defect identification, and studies of wear and tear (Yastikli et al., 2017). In the context of architectural applications, orthophotos are essential for documenting the existing conditions of architectural elements, including facades, roofs, and structural components. By capturing the building as it stands, orthophotos provide a baseline reference for identifying areas in need of repair, deterioration, or modifications Karataş et al., 2022). This capability is particularly valuable in preservation efforts and for managing historical buildings, where maintaining the integrity of the structure is paramount.

Recent advancements in sensor technology have led to substantial improvements in the quality of low-cost sensors. Modern smartphone cameras and compact sensors, often integrated into small unmanned aerial vehicles (UAVs), now offer high-resolution imaging capabilities, improved low-light performance, and enhanced image stabilization features (Elkhrachy, 2021). These sensors have gained popularity due to their low operational costs, ease of use, and versatility, making their applicability for orthophoto production increasingly investigated (Green et al., 2019). For instance, UAVs equipped with such sensors allow for efficient data collection over large areas, enabling rapid assessment and documentation of urban environments and natural landscapes alike (Lahoti et al., 2020).

However, reflective surfaces, such as glass, polished metals, and water bodies, can introduce significant artifacts in the orthophoto generation process that compromise the quality of final product. These challenges arise on the one side because of artifacts like glare and reflections, where bright highlights and distortions can significantly affect the quality and accuracy of the data collected. On the other hand, non-Lambertian surfaces cause significant issues in the traditional photogrammetric workflow.

During the image orientation phase, reflections can cause mismatches between corresponding points in different images, resulting in inaccuracies in tie points, which are crucial for establishing the geometric relationships between images. This issue is exacerbated by the presence of glare, which can distort or eliminate key visual features necessary for accurate feature extraction and matching. Multi-path reflections, wherein light bounces off multiple reflective surfaces before reaching the camera, create ghosting effects that can further confuse the matching algorithms (Karami et al., 2021). These ghosting artifacts not only introduce noise into the captured images but also complicate the dense matching process, which involves computing dense correspondences between image pairs. In such cases, the algorithm may incorrectly associate features from different reflections, leading to erroneous depth information and distortions in the 3D reconstruction (Nicolae et al., 2014). Consequently, the resulting models may fail to accurately represent the real-world scene, particularly in environments with a high prevalence of reflective materials, such as urban settings or industrial sites.

The complexities introduced by reflective surfaces necessitate innovative approaches to mitigate their impact on photogrammetric processes. Researchers have explored various techniques, including adaptive filtering, radiometric normalization, and advanced sensor calibration methods. Adaptive filtering aims to enhance image quality by reducing the effects of glare and reflections, allowing for more accurate feature extraction (Calantropio et al., 2020). Radiometric normalization techniques are employed to adjust the brightness and contrast of images, aiming to compensate for variations in illumination caused by reflections. However, despite these advancements, reflective surfaces remain a significant challenge

in photogrammetry, often leading to inaccuracies that complicate data analysis and interpretation.

The persistent difficulties in accurately reconstructing scenes with reflective surfaces have driven further research into alternative methodologies. For example, recent studies have investigated the integration of machine learning techniques, which can help to identify and classify reflective surfaces more effectively, potentially improving the accuracy of feature matching (Ma et al., 2014). Additionally, emerging approaches such as NeRF and Gaussian splatting have shown promise in enhancing the robustness of reflective surface generation by effectively managing reflections and providing more visually appealing renderings (Croce et al., 2024).

This paper investigates the orthophoto generation process by integrating Gaussian splatting into the workflow. Gaussian splatting is a novel technique that involves projecting Gaussian functions onto a 2D image plane to represent 3D points effectively. Each point in 3D space is represented as a Gaussian blob or "splat," which inherently encapsulates information about its location, size, and shape. This representation allows for the distribution of the point's influence over a region of the image, leading to smooth and continuous renderings that accurately capture the spatial distribution of the underlying data. The inherent properties of Gaussian splats allow for smoother transitions and gradients in the rendered images, which can significantly enhance the quality of orthophotos produced in challenging environments. These factors collectively suggest that Gaussian splatting may represent a promising alternative to produce orthophotos, particularly in the context of low-cost sensors and reflective surfaces. In this paper, we are evaluating the feasibility of Gaussian splats for orthophoto generation.

2. Gaussian splats overview

Gaussian splatting has emerged in the fields of computer graphics, photogrammetry, and remote sensing, effectively addressing the challenges associated with rendering and visualizing 3D data. Unlike traditional rendering methods that rely on meshes or point clouds, Gaussian splatting represents 3D points as Gaussian distributions, or "blobs." Each point in a 3D point cloud is associated with a Gaussian function defined by its mean, which indicates the location of the point, and its covariance, which determines the spread or "blur" of the Gaussian. Mathematically, this can be expressed as:

$$G(x,y,z) = A \cdot exp\left(-\frac{(x-\mu_x)^2+(y-\mu_y)^2+(z-\mu_z)^2}{2\sigma^2}\right) \quad (1)$$

where:

A represents the amplitude of the Gaussian,
μ_x,μ_y,μ_z denote the coordinates of the center of the splats
σ controls the Gaussian's spread

This approach allows the influence of each point to be distributed over a region of space, leading to a smooth representation that captures local variations in the data (Kerbl et al., 2023).

The rendering process in Gaussian splatting consists of several key steps. Initially, the 3D points are projected onto a 2D image plane using the camera's intrinsic and extrinsic parameters. Following this projection, the contribution of each Gaussian to the pixel values is calculated based on its distance from the pixel; points closer to the pixel have a stronger influence. Subsequently, the contributions from multiple overlapping Gaussians are accumulated, resulting in a final pixel color value that represents the combined influence of all contributing Gaussians. To enhance visual quality, post-processing techniques such as tone mapping and filtering can be applied to the rendered image, reducing artifacts and improving the overall appearance (Wu et al., 2024). Efficient computation of the splatting operation is critical, especially for large datasets. Techniques such as the Fast Fourier Transform (FFT) can be employed to speed up the convolution process, as Gaussian functions are particularly amenable to frequency domain manipulation.

Gaussian splatting also incorporates concepts from kernel density estimation (KDE), allowing it to represent point distributions in a more statistically robust manner. Each Gaussian blob can be viewed as a kernel that contributes to a density estimation of the spatial data, which helps in visualizing point clouds with varying densities. This flexibility in handling data distributions makes Gaussian splatting particularly effective in scenarios where traditional methods struggle, such as in the presence of occlusions or varying reflectivity across surfaces. Furthermore, the ability to adjust parameters like bandwidth and spread enables users to fine-tune the output for specific applications, whether for architectural documentation or complex scene reconstruction. By employing techniques like adaptive bandwidth selection, Gaussian splatting can adapt to local data characteristics, enhancing its usability in real-world applications (Chen and Wang, 2024).

The advantages of Gaussian splatting are particularly noteworthy. One significant benefit is its ability to handle sparse and irregular data effectively. This is especially useful in scenarios where point clouds may be unevenly sampled or where data acquisition is inherently challenging, such as in environments with reflective surfaces. Furthermore, the smooth nature of Gaussian functions minimizes the visibility of artifacts that are common in traditional rendering approaches, leading to visually appealing images, with less noticeable artifacts compared to traditional methods that can suffer from issues like aliasing and noise Additionally, Gaussian splatting is computationally efficient, making it suitable for use with low-cost sensors in resource-constrained environments. Its rapid generation of high-quality visualizations provides significant advantages in real-time applications.

Recent studies have demonstrated the potential of Gaussian splatting in various applications, including virtual reality, augmented reality, and environmental monitoring (Li et al., 2025). Its adaptability to different data types and its ability to produce high-quality visual outputs make it an appealing option for researchers and practitioners in the fields of photogrammetry, virtual and augmented realty. As the demand for accurate and reliable 3D models continues to grow across diverse sectors, the integration of Gaussian splatting into orthophoto generation workflows could represent a significant advancement in the field.

A notable comparison arises between Gaussian splatting and Neural Radiance Fields (NeRF). While both techniques are used for 3D representation, they operate on fundamentally different principles. NeRF utilizes deep learning to synthesize novel views of a scene by modelling the volumetric scene as a continuous function. It takes a set of images from different viewpoints and learns to predict the colour and density of points in space, allowing for photorealistic view synthesis (Mildenhall et al., 2021). This approach excels in creating complex lighting effects and can generate images with intricate details and shading, particularly in well-lit environments. However, it often requires substantial computational resources and training data, which can be a limiting factor in real-time applications.

In contrast, Gaussian splatting simplifies the rendering process by directly representing points as Gaussian distributions. This method allows for efficient processing of sparse and irregular datasets without the need for extensive training or computation, making it more suitable for real-time applications and scenarios with limited resources. Gaussian splatting focuses on the spatial

distribution of data rather than learning a complex model, which can be advantageous in certain environments, especially those involving reflective surfaces or incomplete data.

3. Materials and method

3.1 Orthophoto workflows

In this paper, we aim to test and compare two different workflows for orthophoto generation: the conventional photogrammetric method and a Gaussian splatting-based approach. An overview of both methodologies is provided in Figure 1, illustrating the key stages of each process.

Both workflows typically begin with image acquisition, where UAVs or ground-based cameras capture images of the target area or object from multiple viewpoints to ensure comprehensive coverage. This stage is critical as the quality and consistency of the images directly affect subsequent steps. However, when dealing with reflective or transparent materials, such as glass or metal surfaces, specific challenges arise. These materials can distort imagery due to specular reflections and refraction, which alters the appearance of the object and complicates the reconstruction process. Reflections from sunlight or artificial light sources often create glare, obscuring key features and causing difficulties in accurately identifying and matching points between images. Once the images are captured, both workflows require image alignment and orientation, a task typically achieved through a structure-from-motion (SfM) strategy. SfM operates by extracting key points from overlapping images and using these points to reconstruct both the camera's position and orientation. In addition, it generates a sparse 3D point cloud composed of tie points—the corresponding features identified across multiple images. For well-textured, opaque surfaces, this process generally works efficiently, with the algorithm reliably matching features between images. However, for reflective or transparent surfaces, specular reflections and refraction can introduce significant challenges. Reflected light may result in mismatches between corresponding points, as the visual features captured in different images may vary based on the angles of reflection or refraction. This can lead to erroneous or insufficient tie points, significantly affecting the accuracy of the overall orientation. In cases where the number of extracted tie points is too low, the compensation may fail to compute an accurate orientation, resulting in a poorly alignment of the images. Next, the workflow involves georeferencing where Ground Control Points (GCPs) or other forms of real-world data, such as camera positions computed with Post-Processed Kinematic (PPK) techniques, are used to align the model within a geographic coordinate system. This step ensures that the 3D model and the resulting orthophoto is spatially accurate and can be used in applications like mapping or site analysis.

At this stage, the two workflows begin to diverge. In the conventional photogrammetric workflow. Multi-view stereo (MVS) is used to generate a dense 3D point cloud by estimating depth information from overlapping images. While MVS typically excels at reconstructing well-textured surfaces, reflective and transparent materials pose significant challenges. Reflections can confuse the MVS algorithms, as light bouncing off different surfaces may be misinterpreted as depth information. This can result in distorted or incomplete reconstructions, particularly in areas where reflections dominate the captured images. Similarly, transparent objects, such as glass, may refract light in ways that lead to incorrect depth calculations, further complicating the model. Finally, the actual orthophoto generation takes place, where the images are projected onto the 3D model to produce a distortion-free view associated with georeferencing information such as pixel size of the orthophoto and its location. The accuracy of the orthophoto largely depends on the quality of the 3D reconstruction. For standard, opaque surfaces, this process yields high-resolution, geometrically accurate images. However, for reflective or transparent objects, the artifacts and distortions introduced in earlier stages often result in poor-quality orthophotos. These artifacts manifest as ghosting, blurring, or gaps in the orthophoto, particularly in regions where reflections dominate the image. To create orthophotos using the conventional photogrammetric approach, we employed Metashape, a widely used software known for its robust structure-from-motion (SfM) and multi-view stereo (MVS) capabilities. Metashape enabled us to orient images, generate 3D meshes and create orthophotos.

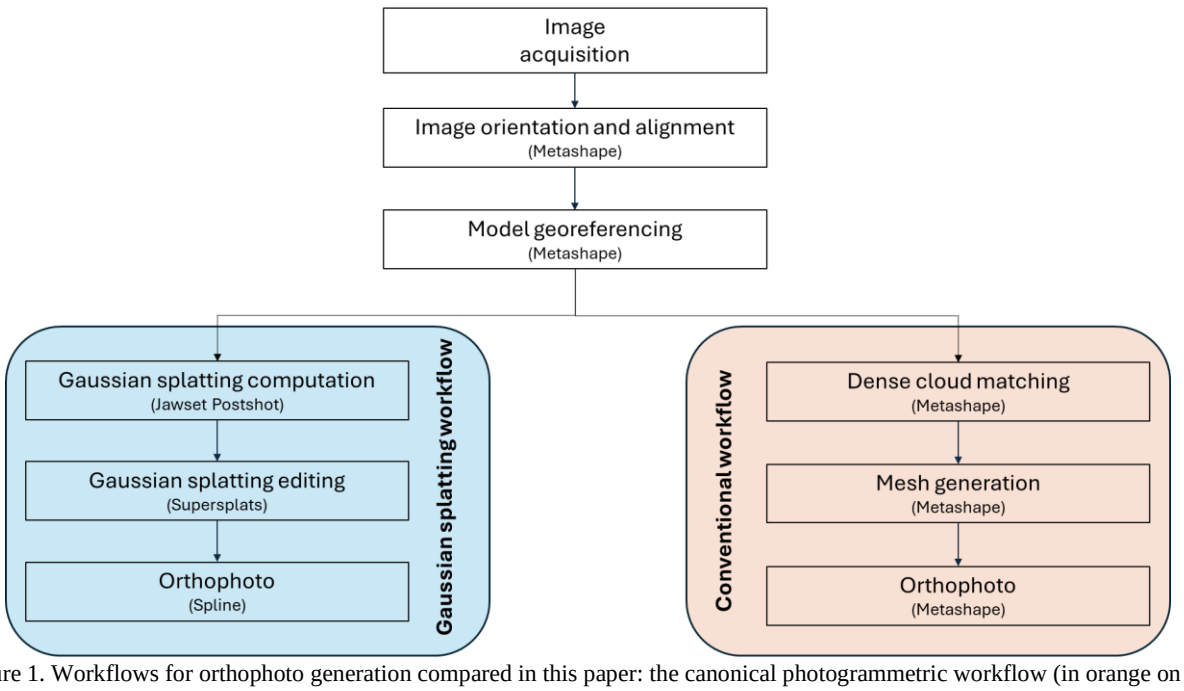

Figure 1. Workflows for orthophoto generation compared in this paper: the canonical photogrammetric workflow (in orange on the right) and the proposed Gaussian splatting workflow (in blue on the left).

In parallel, we evaluated the Gaussian splatting workflows. The first step in the Gaussian splatting workflow involves using Jawset PostShot (https://www.jawset.com/), a software tool designed for advanced 3D rendering based on Gaussian functions. The process begins with importing the image orientation parameters generated by Metashape. This choice was made to ensure consistency with the orientation used in the conventional photogrammetric workflow, rather than computing the image orientation directly in PostShot using the COLMAP implementation available within the software package. Importing parameters from Metashape allows us to maintain the alignment and spatial relationships established during the initial processing stages. At this stage, it appears that adding control points directly in PostShot is not feasible. This limitation emphasizes the importance of having a reliable orientation framework in place before transitioning to the Gaussian splatting workflow. After importing the orientation parameters, it is important to set the images to their original size within PostShot. This is a critical step because altering the image dimensions can lead to changes in the calibration parameters, ultimately affecting the accuracy of the Gaussian splatting process. Keeping the original size preserves the spatial fidelity of the captured data and prevents the introduction of geometric distortions that could compromise the final output. At this point the Gaussian splats are computed using the sparse point cloud data obtained as result of the orientation phase. Each 3D point is represented as a Gaussian blob, allowing for smoother visual representations. In this stage, the software is configured to use all images for training and perform 30,000 iterations, a parameter that plays a significant role in determining the quality of the rendering. Higher iteration counts typically lead to more refined and detailed visual outputs, as they allow the algorithm to converge on an optimal representation of the 3D data. Currently, PostShot supports two distinct splat profiles, each designed to optimize the rendering process for different scenarios:

- Splat MCMC Profile: This profile employs a Markov Chain Monte Carlo (MCMC) approach, which allows users to limit the number of splat primitives generated. By doing so, it effectively reduces the memory and disk space requirements of the resulting model. This is particularly advantageous when working with large datasets or in resource-constrained environments, as it streamlines the rendering process without significantly compromising visual fidelity.
- Splat ADC Profile: Similar to the Splat MCMC profile, the Splat ADC (Adaptive Density Control) profile also controls the number of primitives used. However, it differs in its method of producing scene detail during the training phase. The ADC profile focuses on adapting the density of the splats based on the complexity of the scene, allowing for more nuanced detail in areas requiring higher fidelity. This approach can enhance the quality of renderings in scenarios where fine details are essential, such as reflective or transparent surfaces.

In our test we used Splat MCMC Profile with maximum splats count to 5,000.

Once the training phase of Gaussian splatting is complete, the resulting Gaussian model is exported as a PLY file. This format is widely used for storing 3D Gaussian splatting models. The exported PLY file can then be imported into Supersplat. (https://github.com/playcanvas/supersplat). In Supersplat, users have the capability to refine their 3D models by performing various editing tasks. One of the key features of this software is its ability to identify and remove artifacts, such as ghosting effects. This tool allows for precise manipulation of splat properties, ensuring a cleaner and more accurate representation of the scene. Once the necessary edits are made in Supersplats, the refined model is imported into Spline (https://spline.design/), a platform specialized in 3D modelling and rendering. Within Spline, users can conduct further edits, including cropping the area of interest, adjusting parameters such as splat size, position, and opacity to enhance the final output, contributing to a more realistic rendering. A crucial step in preparing for orthophoto generation is defining the orthogonal plane. This plane serves as the reference for creating orthographic renderings, ensuring that the final output maintains geometric accuracy. Finally, once the orthogonal plane is established, Spline can generate the orthographic rendering of the scene.

3.2 Test datasets

The proposed methodology was tested on several datasets captured in different scenarios like, metallic mechanical elements, urban environments with a prevalence of reflective surfaces and compared with the traditional photogrammetric workflow.

In particular, three main datasets were used (Figure 2):

- “San Michele” dataset;
- “Industrial and transparent” dataset; and
- “Architectural reflective” dataset.

A summary of the datasets is presented in Table 1.

Dataset	Typology	Number of images	Camera type and image size	Description
San Michele	UAV dataset	224	DJI Mavic 3E 5280×3956	Cultural Heritage, Good texture, Complex shape
Industrial and transparent (Synthetic_Metallic)	Synthetic close-range dataset	300	Virtual pinhole camera, 1080x1920 px	Textureless, Complex, Reflective,
Industrial and transparent (Bottle)	Close-range dataset	300	Huawei p20 pro, 1080x1920 px	Complex shape, Highly refractive
Architectural Glass Case	Architectural reflective	140	Samsung A32 2604x4624 px	Architectural, highly reflective, glass surfaces

Table 1. Test dataset overview and data.

a.  b.  c.  d. 

Figure 2. Test dataset overview: “San Michele” (a), “Synthetic_Metallic” (b), “Bottle” (c) and “Architectural Glass Case”.

The first test dataset, referred to as "San Michele", was a feasibility study conducted using a typical UAV dataset aimed at surveying the church of San Michele. This dataset is particularly challenging due to the presence of complex architectural features, such as the bell tower, along with the surrounding buildings within the hamlet. Furthermore, the nearby presence of trees introduces an additional challenge, as vegetation can complicate the generation of high-quality orthophotos due to irregular textures and occlusions. Despite these complexities, the dataset presents well-textured surfaces that are expected to pose minimal issues for both the conventional photogrammetric workflow and the Gaussian splatting method. The objective of the San Michele test is to assess the viability of the proposed approach, ensuring that orthophoto with Gaussian splatting can effectively handle diverse and complex structures, as well as natural obstructions, in comparison to traditional methods. The "San Michele" complex was surveyed with 224 UAV images. The raw data of the GNSS receivers of the drone were used to compute in PPK mode the camera positions. As a base station a virtual master station was computed in the areas using the network of permanent GNSS stations SPIN3 GNSS (https://www.spingnss.it/).

The second dataset, referred to as "Industrial and Transparent", was constructed using data from the publicly available NeRFBK dataset (Yan et al., 2023), accessible on GitHub (https://github.com/3DOM-FBK/NeRFBK). This dataset offers a comprehensive collection of images and 3D data, specifically curated for the evaluation of view synthesis methods and 3D reconstruction techniques. The NeRFBK dataset is particularly valuable for testing in challenging environments due to its diverse range of materials, including transparent and reflective surfaces—factors that often pose significant difficulties in traditional photogrammetric workflows. The dataset's inclusion of industrial scenes, with their complex geometry and reflective materials, makes it ideal for assessing the capabilities of both canonical photogrammetric techniques and the proposed Gaussian splatting approach. This allows for rigorous testing of reconstruction accuracy and the handling of difficult materials like glass and metal. For our "Industrial and Transparent" dataset, we utilized two subsets: "Synthetic_Metallic" and "Bottle." The "Synthetic_Metallic" dataset is a synthetic collection featuring a metal mechanical part with complex geometry, minimal texture, and a highly reflective surface. This dataset was generated using a virtual pinhole camera, capturing 300 images at a resolution of 1080x1920 pixels. Its poor texture and reflective properties pose challenges for traditional photogrammetric techniques. The second dataset, "Bottle," consists of 300 real images captured using a Huawei P20 Pro smartphone, documenting a transparent glass bottle from various angles. The primary processing challenge for this dataset lies in image matching, as refraction and specular reflections caused by light passing through the bottle surface significantly complicate feature detection and matching. Both datasets test the robustness of Gaussian splatting and photogrammetric methods when dealing with challenging material properties such as reflectivity and transparency. For this dataset, the ground truth data was used to establish the scale and align the model within a consistent reference system. This alignment also provided a basis for evaluating the performance of the two reconstruction techniques and orthophoto generation workflows.

		Conventional workflow		**Gaussian splatting workflow**	
Dataset	**Orientation and georeferencing**	**Dense Reconstruction and mesh generation**	**Orthophoto**	**Gaussian spalts generation**	**Orthophoto**
San Michele	Achieved	No evident artifacts or gaps	No artifacts on the buildings some artifacts on vegetation	No evident artifacts or gaps	No artifacts on buildings, overall resolution bit lower than the one of conventional workflow
Industrial and transparent (Synthetic_Metallic)	Achieved	Evident artifacts and missing parts	Evident artifacts	No evident artifacts or gaps	No evident artifacts or gaps, surface smoother than original photos
Industrial and transparent (Bottle)	Achieved	Not achieved	Not achieved	Some artifacts and ghosts effects	Some artifacts and ghosts effects
Architectural Glass Case	Achieved	Evident artifacts and missing parts	Evident artifacts	Some artifacts on the roof of the case	No evident artifacts or gaps

Table 2. Results overview for the different datasets.

The third and final dataset, referred to as the "Architectural Glass Case," focuses on the survey of a structure primarily composed of glass surfaces. This dataset presents a particularly challenging scenario due to the extensive presence of reflective glass elements, which constitute the majority of the building’s exterior. The reflective nature of these surfaces captures and mirrors nearby objects, such as surrounding buildings and trees. Additionally, while there are some opaque surfaces within the

scene, these are entirely textureless, further increasing the difficulty of reconstruction.The primary aim of this dataset is to assess the feasibility of using Gaussian splatting to generate orthophotos for complex, reflective architectural structures. The dataset consists of 140 images captured from various angles around the glass structure using a smartphone. To establish an accurate scale and georeference the model, six points measured with a laser scanner were incorporated.

Figure 3. Results of orthophoto derived from gaussian splatting models: a) “San Michele” dataset (from left to right: image of the dataset, orientation results, gaussian splatting orthophoto results); b) “Synthetic Metallic” dataset (from left to right: image of the dataset, orientation results, orthophoto form standard approach, and orthophoto from gaussian splatting approach); c) “Bottle” dataset (from left to right: image of the dataset, orientation results, dense clods reconstruction results, gaussian splatting orthophoto results); and d). “Architectural Glass Case” dataset (from left to right: image of the dataset, orientation results, orthophoto form standard approach, and orthophoto from gaussian splatting approach)

4. Results

Results on the different datasets are summarised in Table 2 and Figure 3. In comparing the conventional workflow for orthophoto generation with the Gaussian splatting workflow, we can observe notable differences in outcomes across various datasets.

For the San Michele dataset, which includes the complex architecture of the church along with surrounding vegetation, the conventional workflow demonstrated effectiveness in generating a dense mesh without evident artifacts. The overall resolution was maintained, providing a reliable reconstruction of the scene. However, minor artifacts were noted in areas with vegetation, likely due to the challenges posed by shadows and reflections from the foliage. The Gaussian splatting workflow demonstrated comparable completeness but showed slightly lower resolution in the resulting orthophoto compared to that produced by

conventional workflows. This suggests that while the Gaussian splatting method is effective in generating a complete and reliable orthophoto, the conventional workflow remains superior for this type of application. Moving on to the Industrial and Transparent dataset, specifically the Synthetic Metallic component, the conventional workflow produced evident artifacts and gaps in the reconstruction due to the highly reflective nature of the metallic surface. These artifacts can distort the final orthophoto in a significant way determining evident inaccuracies (Figure 4a). Conversely, the Gaussian splatting workflow managed to achieve a smoother surface without these evident artifacts. This success can be attributed to the method's ability to handle the complexities of reflective surfaces more adeptly. The Bottle dataset, which presents unique challenges due to its transparent nature. The conventional workflow struggled significantly, producing artifacts and ghosting effects, particularly around the areas of refraction where light passes through the glass. This workflow failed in generating a reliable model and an orthophoto. In contrast, the Gaussian splatting method not only avoided these ghosting effects but also delivered a more refined output overall. The approach's ability to represent complex light interactions, such as those found in transparent objects, showcases its potential advantages in situations where traditional methods fail. The results from the "Architectural Glass Case" dataset demonstrate promising outcomes for both the glass surfaces and the textureless areas. In the glass panels, the Gaussian splats method successfully captures details such as the glass stickers, which are entirely absent in traditional mesh-based reconstructions. Similarly, the Gaussian splats orthophoto provides a more complete representation of the textureless metallic panels, which are only partially reconstructed in the conventional workflow. By using Gaussian splats instead of traditional mesh-based methods, this approach effectively captures areas with fewer distinct features, like smooth glass panels, by applying a smoothing effect that reduces inconsistencies commonly encountered in photogrammetric workflows. This smoothing also mitigates issues related to glare and reflections, offering a clearer and more accurate visualization of reflective and low-texture surfaces.

Overall, these observations highlight the strengths and weaknesses of each approach in dealing with reflective and transparent materials. The conventional workflow demonstrates reliability in standard conditions but can struggle significantly when faced with complex surfaces that introduce light distortion. In contrast, the Gaussian splatting workflow shows great promise for improving orthophoto generation in challenging scenarios, particularly by providing more reliable outputs in environments where reflections and refractions complicate traditional photogrammetric processes. Initial observations indicate that the Gaussian spatting approach can be an effective technique for addressing the challenges posed by reflective surfaces in orthophoto generation. This was particularly evident in areas with complex reflective structures, where traditional methods would typically fail to produce usable orthophotos. The splatting process also preserved critical details in non-reflective regions, maintaining the overall sharpness and accuracy of the orthophotos without compromising the geometric accuracy of the final product. Moreover, the approach proved to be computationally efficient, making it suitable for use with low-cost sensors that are often deployed in resource-constrained environments. The improved radiometric consistency and artifact reduction observed in the orthophotos suggest that the proposed methodology can significantly enhance the quality of orthophotos generated using low-cost sensors, broadening their applicability in various fields.

One significant limitation of using Gaussian splatting for orthophoto generation is that the resulting orthophoto is not inherently georeferenced. This means that critical metadata files, such as a .tfw (text file), which typically provide information on the geographic position and scale of the orthophoto, are not generated by default. Although it is possible within the Gaussian splatting workflow to define the area of the image and set the desired resolution, thereby determining the pixel size, this only addresses the spatial scale of the orthophoto and not its absolute position in a coordinate system. This limitation presents a significant drawback for applications where georeferencing is crucial, where spatial accuracy is paramount for aligning the orthophoto with other geographic datasets.

In practical terms, this limitation means that while the orthophoto produced by Gaussian splatting may have an accurate representation of size and scale, its position relative to real-world coordinates remains undefined unless additional post-processing steps are taken where users would have to manually georeference the orthophoto in a separate software.

a. b.

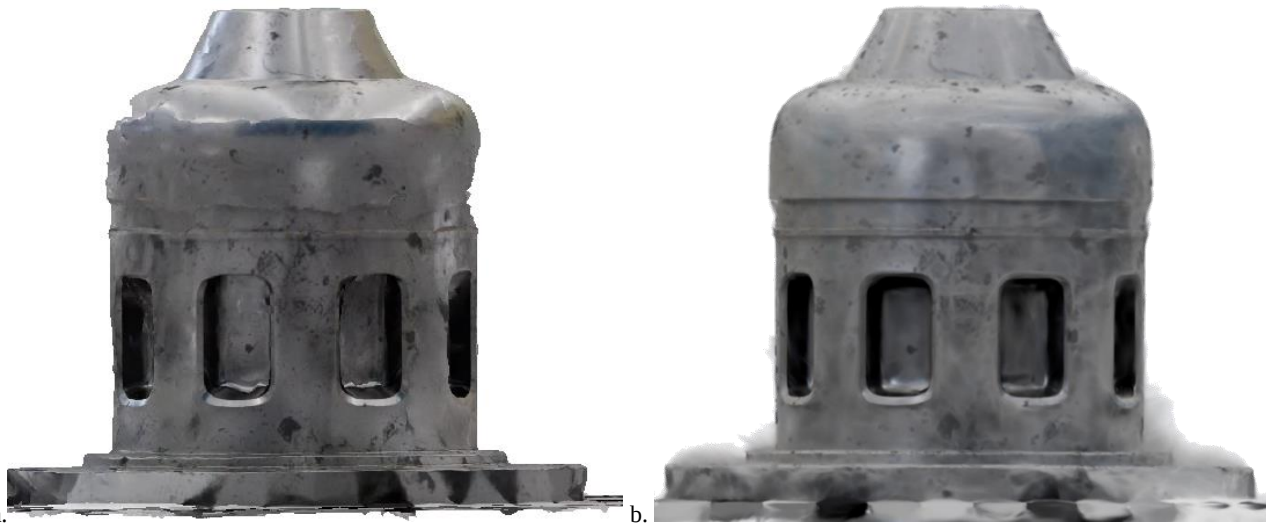

Figure 4. Comparison of orthophotos for the "mechanical" datasets: a) "standard" approach results (see artefact in the top part of the element) and "gaussian splatting result".

5. Conclusions

While the proposed methodology of Gaussian splatting shows great promise, there are notable limitations that warrant careful consideration. One significant challenge is the reliance on accurately identifying reflective surfaces, which can be particularly complex in environments with varying lighting conditions and intricate geometries. This difficulty can lead to inconsistencies in feature matching and ultimately impact the quality of the generated orthophotos. Moreover, the smoothing effect inherent in the splatting process, while effective in minimizing artifacts, can inadvertently obscure fine details in critical areas, resulting in a loss of information that is vital for applications requiring high spatial resolution, such as architectural documentation and structural analysis between artifact reduction and detail preservation highlights the necessity for a balanced approach in implementing Gaussian splatting, especially in contexts where accuracy is paramount. As such, future work will focus not only on refining the splatting process but also on exploring its application to diverse types of imagery, including thermal and multispectral data. Integrating advanced machine learning techniques, such as convolutional neural networks (CNNs) and deep learning frameworks, can potentially enhance the effectiveness of Gaussian splatting by improving the classification of reflective surfaces and optimizing the blending of splats in varying conditions.
Furthermore, develop robust evaluation metrics that can quantitatively assess the quality of orthophotos produced using this methodology, enabling to fine-tune parameters based on specific application requirements.

References

Calantropio, A., Chiabrando, F., Seymour, B., Kovacs, E., Lo, E., Rissolo, D., 2020: Image pre-processing strategies for enhancing photogrammetric 3D reconstruction of underwater shipwreck datasets. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43, 941-948.

Chen, G., Wang, W, 2024: A survey on 3D Gaussian splatting. arXiv preprint arXiv:2401.03890.

Croce, V., Billi, D., Caroti, G., Piemonte, A., De Luca, L., Véron, P., 2024: Comparative Assessment of Neural Radiance Fields and Photogrammetry in Digital Heritage: Impact of Varying Image Conditions on 3D Reconstruction. *Remote Sensing*, 16(2), 301.

Elkhrachy, I., 2021. Accuracy assessment of low-cost unmanned aerial vehicle (UAV) photogrammetry. *Alexandria Engineering Journal*, 60(6), 5579-5590.

Green, D. R., Hagon, J. J., Gómez, C., Gregory, B. J., 2019: Using low-cost UAVs for environmental monitoring, mapping, and modelling: Examples from the coastal zone. In *Coastal management*, 465-501.

Karami, A., Menna, F., Remondino, F., 2021: Investigating 3D reconstruction of non-collaborative surfaces through photogrammetry and photometric stereo. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43, 519-526.

Karataş, L., Alptekin, A., Yakar, M., 2022: Creating Architectural Surveys of Traditional Buildings with the Help of Terrestrial Laser Scanning Method (TLS) and Orthophotos: Historical Diyarbakır Sur Mansion. *Advanced LiDAR*, 2(2), 54-63.

Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023: 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4), 139-1.

Lahoti, S., Lahoti, A., Saito, O., 2020. Application of unmanned aerial vehicle (UAV) for urban green space mapping in urbanizing Indian cities. *Unmanned Aerial Vehicle: Applications in Agriculture and Environment*, 177-188.

Li, Y., Lyu, C., Di, Y., Zhai, G., Lee, G. H., Tombari, F, 2025: Geogaussian: Geometry-aware gaussian splatting for scene rendering. *In European Conference on Computer Vision*, 441-457.

Ma, J., Jiang, X., Fan, A., Jiang, J., Yan, J., 2021: Image matching from handcrafted to deep features: A survey. *International Journal of Computer Vision*, 129(1), 23-79.

Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R., 2021: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), 99-106.

Nicolae, C., Nocerino, E., Menna, F., Remondino, F., 2014: Photogrammetry applied to problematic artefacts. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 40, 451-456.

Yan, Z., Mazzacca, G., Rigon, S., Farella, E. M., Trybala, P., Remondino, F., 2023: NeRFBK: a holistic dataset for benchmarking NeRF-based 3D reconstruction. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48(1), 219-226.

Yastikli, N., Erisir, Z., Altintas, P., Cak, T., 2017: The Use of the Terrestrial Photogrammetry in Reverse Engineering Applications. *In 3D Printing: Breakthroughs in Research and Practice*, 241-250.

Wu, T., Yuan, Y. J., Zhang, L. X., Yang, J., Cao, Y. P., Yan, L. Q., Gao, L., 2024: Recent advances in 3d gaussian splatting. *Computational Visual Media*, 10(4), 613-642.